

REMARKS

This Amendment is filed in response to the Office Action mailed on April 17, 2008, and is herewith filed a Request for Continuing Examination. All objections and rejections are respectfully traversed.

Claims 1-41 are currently pending.

Request for Interview

The Applicant respectfully requests a telephonic interview with the Examiner after the Examiner has had an opportunity to consider this Amendment, but before the issuance of the next Office Action. The Applicant may be reached at 617-951-3067.

Claims Rejections - 35 USC § 103

At paragraphs 2-3 of the Office Action, claims 19 and 21 were rejected under 35 U.S.C. § 103 as being unpatentable over Dada, U.S. Patent Application Publication No. 2004/0093347 A1, hereinafter Dada.

The present invention, as set forth in representative claim 19, comprises in part:

19. A computer readable medium, including program instructions executing on a computer, the program instructions including instructions for performing the steps of:

(a) for each entry in a first data set, placing the entry in a hash table, wherein the first data set is stored on a source storage system;

(b) selecting an entry from a second data set, wherein the second data set is located on a destination storage system and the source storage

system and the destination storage system are separate stand alone storage systems;

- (c) looking up the selected entry in the hash table;
- (d) ***removing, in response to locating the selected entry in the hash table, the selected entry from the hash table;***
- (e) determining if additional second data set entries exist;
- (f) looping to step (b) in response to identifying additional second data set entries; and
- (g) ***reporting a difference between the first data set and the second data set in response to at least one first data set entry remaining in the hash table.***

By way of background, Dada discloses a mechanism for comparing content within data structures. (Dada, Abstract). Specifically, Dada discloses XML documents as the data structures, wherein the content to be compared is marked by XML tags within the documents. (Paragraphs 0016-0017). In Dada, each XML document has a separate hash table. (Paragraphs 0022, 0023). XML tags extracted from each document serve as keys to the document's corresponding hash table. (Paragraphs 0022, 0023). As noted by the Examiner, Dada does not show "the hash values in the *same* hash table" as taught by Applicant. (First Office Action, Paragraph 7) (emphasis added).

Significantly, to compare content, Dada compares every key of a document to every key of every other document. (Dada, Paragraphs 0026, 0028). If a key from one document is missing from the hash table of another document, Dada specifies a difference between the documents. (Paragraph 0026). If no keys are missing, Dada then compares the values associated with the keys. (Paragraph 0027).

In sharp contrast, Applicant claims "***removing, in response to locating the selected entry in the hash table, the selected entry from the hash table***".

Applicant's claimed invention yields significant improvements over brute force comparison techniques, such as Dada, which require an exponential increase in time for each additional element to be compared. (See Applicant's Background, page 7, line 24 – page 8, line 6). Instead of comparing every key of a document to every key of every other document, Applicant claims “*removing, in response to locating the selected entry in the hash table, the selected entry from the hash table*”. Advantageously, as Applicant's procedure progresses, hash entries are removed, thereby enabling lookups to occur more and more quickly. (Applicant's Specification, page 9, lines 16-18).

Nowhere does Dada teach or suggest “*removing, in response to locating the selected entry in the hash table, the selected entry from the hash table*”. In reference to this element of Applicant's claims, the Examiner cited Paragraph 0025 of Dada. (Office Action, Paragraph 7). Paragraph 0025 of Dada, however, discloses creating “a superset of keys that includes all of the keys appearing in either hash table.” (Dada, Paragraph 0025). Paragraph 0025 does not disclose removing any entries from any hash table, therefore Applicant respectfully urges that the cited paragraph in no way renders obvious Applicant's “*removing, in response to locating the selected entry in the hash table, the selected entry from the hash table*”, as claimed.

Applicant respectfully urges that Dada is legally insufficient to render Applicant's claims obvious under 35 U.S.C. § 103(a) because of the absence from Dada of Applicant's claimed novel *removing, in response to locating the selected entry in the hash table, the selected entry from the hash table* and *reporting a difference between the first*

data set and the second data set in response to at least one first data set entry remaining in the hash table.

At paragraphs 4 of the Office Action, claims 1-3, 5-6, 9, 17, 20, 22-26, and 39-40 were rejected under 35 U.S.C. § 103 as being unpatentable over Dada, in view of Dermaine et al., “Adaptive set intersections, unions, and differences”, Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 743-752, year 2000, hereinafter Dermaine.

The present invention, as set forth in representative claim 1, comprises in part:

1. A method for comparing a first order-independent data set comprising unique elements with a second order-independent data set comprising unique elements, comprising:

- (a) for each entry in the first data set, placing the entry in a hash table, wherein the first data set is stored on a source storage system;
- (b) *selecting an entry from the second data set, wherein the second data set is located on a destination storage system and the source storage system and the destination storage system are separate standalone storage systems;*
- (c) looking up the selected entry in the hash table;
- (d) *removing, in response to locating the selected entry in the hash table, the selected entry from the hash table;*
- (e) determining if additional second data set entries exist;
- (f) looping to step (b) in response to identifying additional second data set entries; and
- (g) *reporting a difference between the first data set and the second data set in response to at least one first data set entry remaining in the hash table.*

Dermaine discloses finding the intersection, union, or difference of a collection of

sorted sets. Specifically, in difference proofs, elements in common are withheld from the answer. All other elements are reported in the answer.

Applicant respectfully urges that Dada and Dermaine, taken alone or in combination, do not suggest or teach Applicant's claimed novel *selecting an entry from the second data set, wherein the second data set is located on a destination storage system and the source storage system and the destination storage system are separate stand alone storage systems, ..., removing, in response to locating the selected entry in the hash table, the selected entry from the hash table, ..., and reporting a difference between the first data set and the second data set in response to at least one first data set entry remaining in the hash table*. In further detail, Applicant's invention is determining the difference between a first data set located on a source storage system and a second data set located on a destination storage system using hash tables. The source storage system and the destination storage system are separate storage systems. Applicant's invention is a system for determining the differences between two data sets to only send a limited amount of data to the source storage system.

There is no disclosure in either Dada or Dermaine that of determining the difference in data sets where the data sets are stored on separate storage systems.

Accordingly, Applicant respectfully urges that Dada and Dermaine, taken alone or in combination, are legally insufficient to make obvious the presently claimed invention under 35 U.S.C. § 103 because of the absence of the Applicant's claimed novel *selecting an entry from the second data set, wherein the second data set is located on a destina-*

tion storage system and the source storage system and the destination storage system are separate stand alone storage systems, ..., removing, in response to locating the selected entry in the hash table, the selected entry from the hash table, ..., and reporting a difference between the first data set and the second data set in response to at least one first data set entry remaining in the hash table.

At paragraphs 5 of the Office Action, claims 4, and 7-8 were rejected under 35 U.S.C. § 103 as being unpatentable over Dada, in view of Demaine, and in further view of Carteau, US Patent No. 6,606,694, hereinafter Carteau.

At paragraphs 6 of the Office Action, claims 10-11, 18, and 27-28 were rejected under 35 U.S.C. § 103 as being unpatentable over Dada, in view of Demaine, and in further view of Aiken, US Patent No. 6,240,409, hereinafter Aiken.

At Paragraph 7 of the Office Action, claims 12-16 and 29-31 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Dada, in view of Demaine, and in further view of Bailey, et al., US Patent No. 6,473,767, hereinafter Bailey.

Applicant respectfully notes that claims 4, 7-8, 10-16, 18, and 27-31 are dependent claims that depend from independent claims believed to be in condition for allowance. Accordingly, claims 4, 7-8, 10-16, 18, and 27-31 are believed to be in condition for allowance.

At Paragraph 8 of the Office Action, claims 32-36, 38 and 41 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Dada, in view of Demaine, and in further view of Bailey.

The present invention, as set forth in representative claim 32, comprises in part:

32. A system for performing a consistency check of a source directory replicated to a destination directory by comparing entries in the source and destination directories, the system comprising:

one or more storage disks configured to store one or more of a group consisting of the source directory and the destination directory; and
a process configured to compare entries in the source directory with entries in the destination directory by walking the source and destination directories only once, whereby utilization of storage subsystems associated with the source and destination directories is limited by only walking each of the source and destination directories once, and further configured to report a difference between the source directory and the destination directory, wherein the source directory is located on a source storage system and the destination directory is located on a destination storage system and the source storage system and the destination storage system are separate stand alone storage systems.

Bailey discloses a system for maintaining a list of data files as a plurality of directory entries. Also for creating anti-file directory entries having an attribute indicating there is no real data corresponding to the entry.

Applicant respectfully urges that Dada, Demaine, and Bailey, taken alone or in any combination, do not teach or suggest Applicant's claimed novel *a process configured to compare entries in the source directory with entries in the destination directory by walking the source and destination directories only once, whereby utilization of storage subsystems associated with the source and destination directories is limited by only walking each of the source and destination directories once, and further config-*

ured to report a difference between the source directory and the destination directory, wherein the source directory is located on a source storage system and the destination directory is located on a destination storage system and the source storage system and the destination storage system are separate stand alone storage systems. In further detail, Applicant's invention is determining the difference between a source directory located on a source storage system and a destination directory located on a destination storage system. The source storage system and the destination storage system are separate storage systems. Applicant's invention is a system for determining the differences between two directories to only send a limited amount of data to the source storage system.

There is no disclosure in either Dada, Bailey, or Dermaine that of determining the difference in directories where the directories are stored on separate storage systems.

Accordingly, Applicant respectfully urges that Dada, Dermaine, and Bailey, taken alone or in any combination, are legally insufficient to make obvious the presently claimed invention under 35 U.S.C. § 103 because of the absence of the Applicant's claimed novel *a process configured to compare entries in the source directory with entries in the destination directory by walking the source and destination directories only once, whereby utilization of storage subsystems associated with the source and destination directories is limited by only walking each of the source and destination directories once, and further configured to report a difference between the source directory and the destination directory, wherein the source directory is located on a source storage system and the destination directory is located on a destination storage system and the*

source storage system and the destination storage system are separate stand alone storage systems.

At paragraphs 4 of the Office Action, claim 37 was rejected under 35 U.S.C. § 103 as being unpatentable over Dada, in view of Bailey.

The present invention, as set forth in representative claim 37, comprises in part:

37. A system for comparing entries in a source directory with entries on a destination directory to ensure consistency of replicated data between the source and destination directories, the system comprising:

a computer associated with at least one of the source and destination directories, the computer comprising a directory comparison process configured to perform a comparison of entries in the source and destination directories by walking each directory once and placing entries in a hash table and further configured to remove matching entries from the hash table, whereby computational cost is reduced for future look up operations in the hash table, wherein the source directory is located on a source storage system and the destination directory is located on a destination storage system and the source storage system and the destination storage system are separate stand alone storage systems.

Applicant respectfully urges that Dada and Bailey, taken alone or in combination, do not suggest or teach Applicant's claimed novel *a computer associated with at least one of the source and destination directories, the computer comprising a directory comparison process configured to perform a comparison of entries in the source and destination directories by walking each directory once and placing entries in a hash table and further configured to remove matching entries from the hash table, whereby computational cost is reduced for future look up operations in the hash table, wherein*

the source directory is located on a source storage system and the destination directory is located on a destination storage system and the source storage system and the destination storage system are separate stand alone storage systems. In further detail, Applicant's invention is determining the difference between a source directory located on a source storage system and a destination directory located on a destination storage system. The source storage system and the destination storage system are separate storage systems. Applicant's invention is a system for determining the differences between two directories to only send a limited amount of data to the source storage system.

There is no disclosure in either Dada, or Bailey that of determining the difference in directories where the directories are stored on separate storage systems.

Accordingly, Applicant respectfully urges that Dada and Bailey, taken alone or in combination, are legally insufficient to make obvious the presently claimed invention under 35 U.S.C. § 103 because of the absence of the Applicant's claimed novel *a computer associated with at least one of the source and destination directories, the computer comprising a directory comparison process configured to perform a comparison of entries in the source and destination directories by walking each directory once and placing entries in a hash table and further configured to remove matching entries from the hash table, whereby computational cost is reduced for future look up operations in the hash table, wherein the source directory is located on a source storage system and the destination directory is located on a destination storage system and the source storage system and the destination storage system are separate stand alone storage systems.*

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable base claims
and likewise in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account
No. 03-1237.

Respectfully submitted,

/Shannen C. Delaney/
Shannen C. Delaney
Reg. No. 51,605
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500